

REMARKS

In response to the Office Action dated September 20, 2007, the Applicants respectfully request reconsideration. To further the prosecution of this application, each of the rejections set forth in the Office Action has been carefully considered and is addressed below. The application is believed to be in condition for allowance.

Initially, the undersigned thanks Examiner Taylor for the courtesies extended in granting and conducting a telephone interview on December 4, 2007. The substance of that interview is summarized herein.

Election/Restriction

As stated in the Office Action, during a prior telephone interview, the Applicants' representative orally elected the invention of Group I (claims 1-26) without traverse. Affirmation of that election is hereby made. Claims 27-47 have been cancelled.

Claim Objections

Claims 2 and 13 were objected to as including a period inserted in the middle of the claim. Those claims have been amended to address the objection.

Claim Rejections Under 35 USC §112

Claims 4-6 were rejected under 35 USC §112, paragraph 2, as purportedly being indefinite for referring to "the hash" which lacked antecedent basis in independent claim 1. In view of an amendment herein that introduces the recitation of a hash in claim 1, that rejection is believed to be overcome.

Claim Rejections Under 35 USC §102

Claims 1-26 are rejected under 35 USC §102 as purportedly being anticipated by Tolia et al. This rejection is respectfully traversed for the reasons set forth below.

Brief Overview of Embodiments of the Invention

Embodiments of the present invention relate to a content addressable storage (CAS) system in which units of data are accessed using a content address that is based at least in part on the content of the unit of data. In such a system, when a content address is received, the physical location of the unit of data in the storage system is determined, and the unit of data is retrieved from that location and returned to the requesting device (e.g., a host computer). (Specification, page 1, lines 18-21). As described in the Background of the specification, the task of determining the physical location of a unit of data may be involved, as unlike other types of storage systems (e.g., block I/O storage systems), the content address conventionally did not provide any information suggesting the physical storage location of the corresponding unit of data. (Specification, page 1, line 22 – page 5, line 10). In view of that, the specification describes a number of improved techniques for accessing content in a CAS system.

In one embodiment of the invention, a technique is employed to create content addresses for units of data such that the content addresses include information that provides an indication of which units of data were written to the storage system proximate in time. (Page 12, lines 14-17). During the telephone interview, the Examiner questioned what was meant by proximate in time, and the undersigned explained that the term is meant to refer to units of data which were written to the storage system around the same time.

In some embodiments, providing content addresses that have some similarity in at least a portion thereof for units of data that are written proximate in time can increase efficiencies in a storage system locating units of data in response to access requests from a host. (Page 12, lines 17-20). For example, in one embodiment, units of data may be stored in files on a file system, and each file may have a file name corresponding, at least in part, to the content address of the corresponding unit of data. (Page 28, lines 13-14). The hashing algorithm used to generate content addresses in a conventional CAS system may yield fairly random values. (Page 28, lines 20-23). As a result, Applicants appreciated that units of data written at approximately the same time may have vastly different content addresses, and as a result, vastly different file names that result in those content units being stored in different locations in the file system hierarchy. (Page

28, lines 24-26). That is, despite the fact that units of data are written in close proximity in time, the units of data may be scattered across the file system due to the randomness of the hashing algorithm which may preclude any locality of reference in the file system. (Page 28, lines 24-28).

Operating system typically maintain a file system map that maps file system locations to physical disk locations. (Page 28, lines 29-30). When the file system is large, the map maintained by the operating system may also be quite large. (Page 28, lines 30-32). Thus, various portions of the file system map may need to be brought into primary memory and others swapped out as needed. (Page 29, lines 3-5). It should be appreciated that access performance is better when accessing portions of the file system map in memory rather than disk. (Page 29, lines 5-7).

In view of the foregoing, Applicants appreciated that in a storage system that employs a file system structure, it is desirable to store units of data that are stored at approximately the same time in approximately the same file system location to minimize the likelihood of the operating system having to go to disk to load different portions of the file system for sequentially accessed data units. (Page 29, lines 17-21). In some file systems, the file system location of a file is determined based upon the file name. (Page 29, lines 23-24). Thus, in accordance with one embodiment of the invention, information is added to the file name of a file that is similar for files written at approximately the same time. (Page 29, lines 23-25). This provides temporal locality of reference for files written at approximately the same time, and tends to result in those files being stored in proximate file system locations. (Page 29, lines 25-26).

In some embodiments directed to a CAS system, the information that is added to the file name of the file may also be added to the content address of the data stored in the file. (Page 29, lines 27-29). For example, the information may be added to the hash of the content of the file to generate a content address based on the hash and the additional information. (Page 29, lines 29-31).

The above example was described as providing temporal locality of reference with respect to files stored in a file system. However, it should be appreciated that such a technique is

applicable to other logical data storage constructs. For example, adding temporal locality of reference information to content addresses (or other information stored on a database) may increase access efficiency of database tables that are accessed using a content address (or other information) as a key, and can also be used with other logical constructs. (Page 30, lines 23-27).

It should be appreciated that the foregoing overview of embodiments of the invention is provided merely to assist the Examiner in appreciating various aspects of the present invention. However, not all of the description provided above necessarily applies to each of the independent claims pending in the application. Therefore, the Examiner is requested to not rely upon the foregoing summary in interpreting any of the claims or in determining whether they patentably distinguish over the prior art of record, but rather is requested to rely only upon the language of the claims themselves and the arguments specifically related thereto provided below.

Tolia et al.

Tolia et al. is directed to a system that leverages (or makes “opportunistic use of”) content addressable storage to enable the implementation of a distributed file system by employing a concept referred to as “file recipes.” (See title and the abstract). Thus, Tolia employs CAS to improve the performance of a conventional distributed file system. (See page 1, ¶ 2, lines 4-6). The concept of a file recipe is central to Tolia’s approach (see page 1, ¶ 3, line 1) and provides instructions for the organization or construction of a file from a number of data blocks that each is referenced using a block identifier that is a cryptographic hash of the contents of the block. (See page 1, ¶ 3, lines 3-4 and page 2, in the first paragraph under the sub-heading 3 entitled “Recipes”). A file recipe lists the addresses (i.e., the content addresses or hashes) for the data blocks that compose the desired file, describes the arrangement of those component blocks in the file and may also include metadata about the file. (See the first and fifth paragraphs under the sub-heading “3 Recipes”). An example of the metadata about the file can include the time the file was last modified. (See the fifth paragraph under “3 Recipes” at lines 2-5).

Each Of The Claims Patentably Distinguishes Over Tolia

Each of independent claims 1, 12 and 23 is directed to creating content addresses for units of data to comprise first information that provides an indication of which units of data are written to a CAS system proximate in time. In addition, each has been amended to include the further limitation that, for at least one of the units of data, a hash is created for at least a portion of the unit of data and the first information is added to the hash to create the content address for the one of the units of data.

As discussed during the telephone interview, Tolia does not teach creating content addresses for units of data to comprise any information that provides an indication of which units of data were written to the storage system proximate in time. The Office Action (see paragraph 8) asserts that this limitation is disclosed at "Tolia, section 3 recipes discussion." The undersigned has carefully reviewed section 3 of Tolia and the discussion of recipes, and sees nothing disclosed therein about creating any of the content addresses for the units of data to comprise information providing an indication of which were written to the storage system proximate in time. It is noted that the file recipe includes metadata for the file as a whole which specifies a time when the file was last modified. However, as discussed during the interview, that metadata is not included in the content address for any of the data blocks in the file. In addition, that metadata is directed solely to the file as a whole, and might indicate the time that only a portion of the file was last modified. In this respect, it should be appreciated that a file can be modified by modifying only a subset of the data blocks therein, such that an indication of when the file was last modified provides absolutely no indication about which data blocks were modified at that time, and therefore, can provide no indication of which units of data were written to the storage system proximate in time.

During the telephone interview, the Examiners appeared to appreciate these distinctions. They suggested that certain distinguishing aspects of one embodiment of the present invention could be further clarified by introducing into claim 1 limitations similar to those recited in claim 3, relating to creating the content address for a unit of data by creating a hash and adding to the hash the information providing an indication of which units of data were written to the storage system proximately in time. While Applicants reserve the right to continue to pursue claims of

the scope of the original independent claims, each of the independent claims has been amended to include a limitation of this type. It should be appreciated that the language added differs slightly from the language previously pending in the independent claims and refers to “adding” the first information to the hash rather than “inserting” the information in the hash to make clear that the information may be added at any suitable position in the hash (e.g., at the beginning, at the end or at any other position). (See specification, page 29, line 31 – page 30, line 1).

As should be appreciated from the foregoing, Tolia does not teach that any of the content addresses is created by adding to a hash of at least a portion of the unit of data information that provides an indication of which units of data are written to the storage system proximate in time. As such a limitation is included in each of the pending claims, the rejection of each of the pending claims should be withdrawn.

Summary of Other Miscellaneous Issues Discussed During The Telephone Interview

During the telephone interview, some discussion was had concerning whether the independent claims were acceptable in reciting creating a hash address in a novel and non-obvious way, or whether the claims needed to be amended to recite specific uses of the information in the content addresses. The undersigned expressed the view that nothing further was needed, and questioned what bases might exist for requiring additional steps limited to particular uses of the information added to a content address. In this respect, the undersigned explained that there may be numerous ways in which the information can be used, and that one embodiment of the present invention is directed to the creation of such a new and non-obvious content address, such that the Applicants should not be required to limit their claimed invention to a subset of all possible uses.

The Examiner questioned whether claim 1, for example, could be viewed as a single means claim if not amended. The undersigned pointed out that claim 1 is a method claim and does not include any clause reciting “means for” performing any function, let alone being a claim including only such a clause.

The Examiner also pointed out that a claim must have utility. The undersigned explained that the Examiner appeared to be suggesting that the utility for a claimed invention must be explicitly recited in the claim itself, and that there is no such requirement under the Patent Law. Rather, the Patent Law only requires that a claimed invention have utility. The undersigned further pointed out that the inventions recited in the claims all clearly have utility, examples of which are described in the specification.

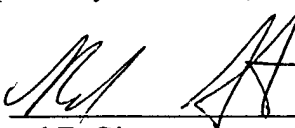
CONCLUSION

A Notice of Allowance is respectfully requested. The Examiner is requested to call the undersigned at the telephone number listed below if this communication does not place the case in condition for allowance.

If this response is not considered timely filed and if a request for an extension of time is otherwise absent, Applicants hereby request any necessary extension of time. If there is a fee occasioned by this response, including an extension fee, that is not covered by an enclosed check, please charge any deficiency to Deposit Account No. 23/2825.

Dated: December 19, 2007

Respectfully submitted,

By: 
Richard F. Giunta
Registration No.: 36,149
WOLF, GREENFIELD & SACKS, P.C.
Federal Reserve Plaza
600 Atlantic Avenue
Boston, Massachusetts 02210-2206
(617) 646-8000